# Recurrent Neural Machine Translation

**Biao Zhang**[1], **Deyi Xiong**[2] **and Jinsong Su**[1]
Xiamen University, Xiamen, China 361005[1]
Soochow University, Suzhou, China 215006[2]
zb@stu.xmu.edu.cn, jssu@xmu.edu.cn
dyxiong@suda.edu.cn

arXiv:1607.08725v1 [cs.CL] 29 Jul 2016

## Abstract

The vanilla attention-based neural machine translation has achieved promising performance because of its capability in leveraging varying-length source annotations. However, this model still suffers from failures in long sentence translation, for its incapability in capturing long-term dependencies. In this paper, we propose a novel recurrent neural machine translation (RNMT), which not only preserves the ability to model varying-length source annotations but also better captures long-term dependencies. Instead of the conventional attention mechanism, RNMT employs a recurrent neural network to extract the context vector, where the target-side previous hidden state serves as its initial state, and the source annotations serve as its inputs. We refer to this new component as *contexter*. As the encoder, contexter and decoder in our model are all derivable recurrent neural networks, our model can still be trained end-to-end on large-scale corpus via stochastic algorithms. Experiments on Chinese-English translation tasks demonstrate the superiority of our model to attention-based neural machine translation, especially on long sentences. Besides, further analysis of the contexter revels that our model can implicitly reflect the alignment to source sentence.

## 1 Introduction

Recently, the neural machine translation (NMT) enhanced with attention mechanism has achieved very promising results on multiple language pairs, including English-French (Bahdanau et al., 2014; Luong et al., 2015b), English-German (Jean et al., 2015; Luong et al., 2015a) and Chinese-English (Shen et al., 2015; Tu et al., 2016; Zhang et al., 2016), etc. Typically, NMT models build upon an *encoder-decoder* framework, where the *encoder* reads and encodes a source language sentence into distributed vectors, from which the *decoder* generates its corresponding target language sentence word by word (Bahdanau et al., 2014; Sutskever et al., 2014). In order to fully leverage the varying-length source annotations according to target-side states, Bahdanau et al. (2014) introduce the attention mechanism, which makes current NMT beyond the early sequence-to-sequence learning (Sutskever et al., 2014).

However, current NMT still suffers from undesirable failures in long sentence translation. The main reason for this lies in that the attention mechanism only models a relative weight for each source annotation separately, but ignores the potential internal dependencies among these annotations, especially the long term dependencies. Consequently, some words, even clauses, are often mistakenly untranslated or over-translated, which is referred to as the *coverage* problem (Tu et al., 2016; Cohn et al., 2016; Feng et al., 2016). Borrowing from conventional SMT, Tu et al. (2016) propose to maintain a *coverage* vector as a guidance to deal with those untranslated and over-translated words and clauses. In spite of their success, their model still suffers from the inability in modeling dependencies inside the source annotations.

In this paper, we propose a recurrent neural machine translation (RNMT) to solve the dependency problem. As recurrent neural network (RNN) can not only deal with varying-length inputs, but also capture long-term dependencies inside these inputs (Hochreiter and Schmidhuber, 1997; Chung et al., 2014), we replace the attention mechanism with a RNN, or more concretely, the GRU model (Chung

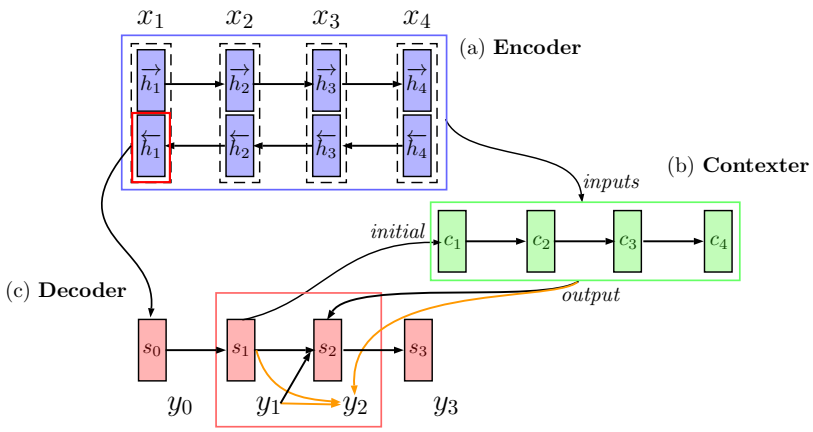


Figure 1: Overview of the proposed recurrent neural machine translation model. We use blue, green and red color to indicate the encoder-related (x), contexter-related (c) and decoder-related (y) representation respectively. The yellow lines show the flow of information employed for target word prediction.

et al., 2014). As the GRU aims at dynamically generating a *context* vector for a target word being generated, we refer to this component as *contexter* in our model. Figure 1 shows the overall architecture. Particularly, we treat the target-side previous hidden state $s_{j-1}$ as the initial state $c_0$ for the contexter, and feed all the hidden states produced by the encoder as its inputs. This enables the output of our contexter to vary according to the target-side translation history. We investigate two different methods for the output of our contexter: *mean-pooling* and *last state*. The former has been verified in various classification problems, while the latter has often been assumed to encode the whole input information. No matter which output, our model is still a unified neural network that can be trained end-to-end via stochastic algorithms.

We compare the proposed model against the vanilla attention mechanism on Chinese-English translation tasks. The empirical results show that our model yields significantly better translation performance in terms of both BLEU and NIST scores, especially on long sentences (with a gain of up to 4.0 BLEU points) where long-term dependencies are heavily required. Besides, we further analyze the gates in GRU and find that our model can reflect the alignment to source sentence implicitly.

Our major contribution lies in the following three aspects:

- We introduce the GRU as an alternate for the attention mechanism, which not only preserves the ability to leverage the varying-length source annotations, but also captures the long-term dependencies.
- We provide a method to visualize the learned GRU in our model. Experiment result reveals that our model can reflect the alignment to source sentence implicitly.
- We conduct experiments on Chinese-English translation tasks. Translation results, especially on long sentences, demonstrate the superiority of our model to the attention mechanism.

## 2 Background

In this section, we briefly review the attention-based NMT which directly estimates the conditional probability $p(\mathbf{y}|\mathbf{x})$ given a collection of source and target sentence pairs $(\mathbf{x}, \mathbf{y})$. We explain the model with regards to the encoder, the decoder and the attention mechanism. The former two are shared in our model, and shown in Figure 1 (a) and (c) respectively. The latter, however, is replaced with GRU in our model. We start from the encoder.

**Encoder.** The encoder is a bidirectional RNN that composes of a forward and a backward RNN. The forward RNN reads the source sentence $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ from left to right while the backward RNN reads in the opposite direction (see the parallel arrows in Figure 1 (a)):

$$\begin{aligned} \overrightarrow{h}_i &= f_{enc}(\overrightarrow{h}_{i-1}, E_{x_i}) \\ \overleftarrow{h}_i &= f_{enc}(\overleftarrow{h}_{i+1}, E_{x_i}) \end{aligned} \tag{1}$$

where $E_{x_i} \in \mathbb{R}^{d_w}$ is the embedding for source word $x_i$, and $\overrightarrow{h}_i, \overleftarrow{h}_i \in \mathbb{R}^{d_h}$ are the hidden states

generated in two directions. Following previous work (Bahdanau et al., 2014), we set the encoding function $f_{enc}(\cdot)$ to a GRU.

The source annotations $\mathbf{H} = (\mathbf{h}_1, \mathbf{h}_2, \ldots, \mathbf{h}_n)$ are obtained by concatenating the forward and backward hidden states:

$$\mathbf{h}_i = \left[\overrightarrow{h}_i^T; \overleftarrow{h}_i^T\right]^T \quad (2)$$

In this way, each annotation $\mathbf{h}_i$ encodes information about the $i$-th word with respect to all the other surrounding words in the source sentence.

**Decoder.** The decoder is a forward RNN, whose initial hidden state $\mathbf{s}_0$ is initialized according to the source-side last hidden state $\overleftarrow{h}_1$. Given the previous target word $y_{j-1}$, the $j$-th hidden state of the decoder $\mathbf{s}_j$ and the context vector $\widetilde{\mathbf{c}}_j$, the decoder calculates the conditional probability of the $j$-th target word as follows (see yellow lines in Figure 1 (c)):

$$p(y_j|y_{<j}, \mathbf{x}) = g(E_{y_{j-1}}, \mathbf{s}_j, \widetilde{\mathbf{c}}_j) \quad (3)$$

where $g(\cdot)$ is a non-linear function, $\widetilde{\mathbf{c}}_j$ is obtained by the attention mechanism and $\mathbf{s}_j$ is calculated as follows:

$$\mathbf{s}_j = f_{dec}(\mathbf{s}_{j-1}, E_{y_{j-1}}, \widetilde{\mathbf{c}}_j) \quad (4)$$

Here, $E_{y_{j-1}} \in \mathbb{R}^{d_w}, \mathbf{s}_{j-1} \in \mathbb{R}^{d_h}$ and $\widetilde{\mathbf{c}}_j \in \mathbb{R}^{2d_h}$. $f_{dec}(\cdot)$ is, again, the GRU. The decoder thus utilizes Eq. (4) and (3) alternatively to generate each word in the target sentence until reaching a predefined end token.

**Attention Mechanism.** The attention mechanism aims at extracting a context vector dynamically according to the previous decoder state. Formally,

$$\widetilde{\mathbf{c}}_j = a(\mathbf{s}_{j-1}, \mathbf{H}) \quad (5)$$

The non-linear function $a(\cdot)$ defines this process. Typically, it first computes a relevance weight $\alpha_{ji}$ of the $i$-th annotation for the $j$-th target word via a feed forward neural network that takes as input $\mathbf{h}_i$ and the previous decoder state $\mathbf{s}_{j-1}$:

$$\alpha_{ji} = \frac{\exp(e_{ji})}{\sum_i \exp(e_{ji})} \quad (6)$$

where the relevance score $e_{ji}$ is computed as follows,

$$e_{ji} = v_a^T \tanh(W_a \mathbf{s}_{j-1} + U_a \mathbf{h}_i) \quad (7)$$

Intuitively, the weights $\{\alpha_{ji}\}_{i=1}^n$ act as the word alignment, and quantify how much each source annotation contributes to the word prediction at each time step. The context vector $\widetilde{\mathbf{c}}_j$ is thus calculated as the weighted summation of $\mathbf{H}$:

$$\widetilde{\mathbf{c}}_j = \sum_i \alpha_{ji} \mathbf{h}_i$$

More details could be found in (Bahdanau et al., 2014).

Because of the ability in modeling source context, attention-based NMT goes beyond the sequence to sequence model. However, as we can see from Eq. (7), the relevance score only models the relation between $\mathbf{s}_{j-1}$ and $\mathbf{h}_i$, which ignores the dependencies among $\{\mathbf{h}_i\}_{i=1}^n$ completely. As a result, attention-based NMT often has poor performance for long sentence translation. We elaborate how to solve this problem in next section.

## 3 Recurrent Neural Machine Translation

The overall architecture of our model is shown in Figure 1, where the encoder and decoder is essentially the same as attention-based NMT. Instead of the attention mechanism, we propose a *contexter* (Figure 1 (b)) as its alternate to extract the context vector. The mathematic form of our contexter is the same as Eq. (5).

We employ the RNN enhanced with GRU to this end because of its capacity in handling long-term dependencies in sequential data. The reason why we do not use the vanilla RNN and the Long Short-Term Memory is that the former often suffers from the vanishing and exploding gradient problem, while the latter is usually computational inefficient.

Concretely, the $t$-th GRU unit in our contexter consists of a *reset* gate $\mathbf{r}_t$, a *update* gate $\mathbf{z}_t$ and a hidden state $\mathbf{c}_t$,

$$\begin{aligned} \mathbf{z}_t &= \sigma(W_z \mathbf{h}_t + U_z \mathbf{c}_{t-1} + b_z) \\ \mathbf{r}_t &= \sigma(W_r \mathbf{h}_t + U_r \mathbf{c}_{t-1} + b_r) \\ \overline{\mathbf{c}}_t &= \tanh(W \mathbf{h}_t + U [\mathbf{r}_t \odot \mathbf{c}_{t-1}] + b) \\ \mathbf{c}_t &= (1 - \mathbf{z}_t) \odot \mathbf{c}_{t-1} + \mathbf{z}_t \odot \overline{\mathbf{c}}_t \end{aligned} \quad (8)$$

where,

$$\mathbf{c}_0 = \tanh(V \mathbf{s}_{j-1} + b_0) \quad (9)$$

$W_* \in \mathbb{R}^{d_h \times 2d_h}, U_* \in \mathbb{R}^{d_h \times d_h}, V \in \mathbb{R}^{d_h \times d_h}$ are weight matrices and $b_* \in \mathbb{R}^{d_h}$ are bias terms. In

this way, our model explicitly integrates the dependencies among source annotations into the hidden state $\mathbf{c}_t$. Intuitively, at each time step for target word prediction, our contexter rereads the source sentence together with current history decoding information, and then extracts important source words for the translation. This, to some extent, is consistent with human behavior.

We explore two approaches to extract the context vector from hidden states of the contexter $\mathcal{C} = (\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_n)$, namely *mean-pooling* and *last state*. The former uses the average of $\mathcal{C}$ as the context vector, which has been demonstrated in various classification tasks:

$$\widetilde{c}_j = \frac{1}{n} \sum_i \mathbf{c}_i \tag{10}$$

The latter uses the last contexter state, i.e. $\widetilde{c}_j = \mathbf{c}_n$, which is often assumed to have encoded all the input information. We conduct experiments to check their effectiveness respectively.

**Model Training** We train the encoder, the decoder and the contexter to maximize the conditional log-likelihood of the training data:

$$\arg\max_{\theta} \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x},\mathbf{y})\in\mathcal{D}} \sum_j \log p_\theta(y_j | y_{<j}, \mathbf{x}) \tag{11}$$

where $\mathcal{D}$ is the bilingual training corpus. As all the components in our model are derivable, we can optimize the model parameters $\theta$ using standard stochastic algorithms.

## 4 Experiments

We conducted a series of experiments on the NIST Chinese-English translation tasks to evaluate the effectiveness of the proposed model.[1]

### 4.1 Setup

**Datasets and Evaluation Metrics** Our training data is a combination of LDC2003E14, LDC2004T07, LDC2005T06, LDC2005T10 and LDC2004T08 (Hong Kong Hansards/Laws/News), which contain 2.9M sentence pairs, with 80.9M Chinese words and 86.4M English words respectively. We used the NIST 2005 (1082 sentences) dataset as our development set, and the NIST 2002 (878 sentences), 2003 (919 sentences), 2004 (1788 sentences), 2006 (1664 sentences), 2008 (1357 sentences) datasets as our test sets. We employed the case-insensitive BLEU-4 metric (Papineni et al., 2002) and the NIST score (Doddington, 2002) to evaluate translation quality. We performed paired bootstrap sampling (Koehn, 2004) for significance test using the script in *Moses*[2].

**Baselines and Settings** We compared our model against two state-of-the-art SMT and NMT systems:

- *Moses* (Koehn et al., 2007): an open source conventional phrase-based SMT system.
- *GroundHog* (Bahdanau et al., 2014): the neural machine translation system enhanced with attention mechanism.

For the *Moses*, we trained a 4-gram language model on the Xinhua section of the English Gigaword corpus (306M words) using the SRILM[3] toolkit with modified Kneser-Ney smoothing. All other parameters were kept as the default settings. Besides, we used all source and target words in the training corpus for this system.

For the *GroundHog*, we set the maximum length of training sentences to be 50 words, and preserved the most frequent 30K words as the source and target vocabulary respectively, which covers 98.9% and 99.2% of the text respectively. All other words were represented by a specific token "*UNK*" as the default implementation. Following Bahdanau et al. (2014), we set $d_w = 620$, $d_h = 1000$. All other settings are the same as the default configuration (for *RNNSearch-50*). We used the Adadelta algorithm for optimization, with a batch size of 80. The model parameters were selected according to the maximum BLEU points on the development set. Besides, during decoding, we used the beam-search algorithm, and set the beam size to 10.

For the *RNMT*, we initialized its parameters with the RNNSearch in *GroundHog*. The settings of our model are the same as that of *GroundHog*, except for the dimension of our contexter which is 1000 in our model, but 2000 in *GroundHog*. We implemented

[1] Source code for the RNMT will be released upon publication.

[2] https://github.com/moses-smt/mosesdecoder/blob/master/scripts/analysis/bootstrap-hypothesis-difference-significance.pl

[3] http://www.speech.sri.com/projects/srilm/download.html

| Metric | System | MT05 | MT02 | MT03 | MT04 | MT06 | MT08 | ALL |
|---|---|---|---|---|---|---|---|---|
| BLEU | *Moses* | 33.68 | 34.19 | 34.39 | 35.34 | 29.20 | 22.94 | 31.80 |
| | *GroundHog* | 31.38 | 33.32 | 32.59 | 35.05 | 29.80 | 22.82 | 31.32 |
| | *RNMT (mean-pooling)* | 32.49 | 34.00 | 33.11 | 36.44$^{+}$ | 30.27 | 23.44$^{+}$ | 32.11$^{\uparrow ++}$ |
| | *RNMT (last state)* | 32.68 | **35.04**$^{++}$ | 33.96$^{++}$ | **37.52**$^{\uparrow ++}$ | **30.97**$^{\Uparrow ++}$ | **24.02**$^{\uparrow ++}$ | **33.11**$^{\Uparrow ++}$ |
| NIST | *Moses* | 9.17 | **9.55** | **9.47** | **9.65** | **8.29** | **6.93** | **9.51** |
| | *GroundHog* | 7.99 | 8.49 | 8.22 | 8.76 | 7.75 | 6.53 | 8.59 |
| | *RNMT (mean-pooling)* | 8.17 | 8.67 | 8.34 | 8.98$^{++}$ | 7.89$^{+}$ | 6.50 | 8.74$^{+}$ |
| | *RNMT (last state)* | 8.22 | 8.85$^{++}$ | 8.50 | 9.21$^{++}$ | 8.00$^{++}$ | 6.75 | 8.95$^{++}$ |

Table 1: BLEU and NIST scores on the NIST Chinese-English translation tasks. **ALL** = score on all test data. We highlight the best results in bold for each test set. "$\uparrow$": significantly better than *Moses* ($p < 0.05$); "$\Uparrow$": significantly better than *Moses* ($p < 0.01$); "+": significantly better than *GroundHog* ($p < 0.05$); "++": significantly better than *GroundHog* ($p < 0.01$);

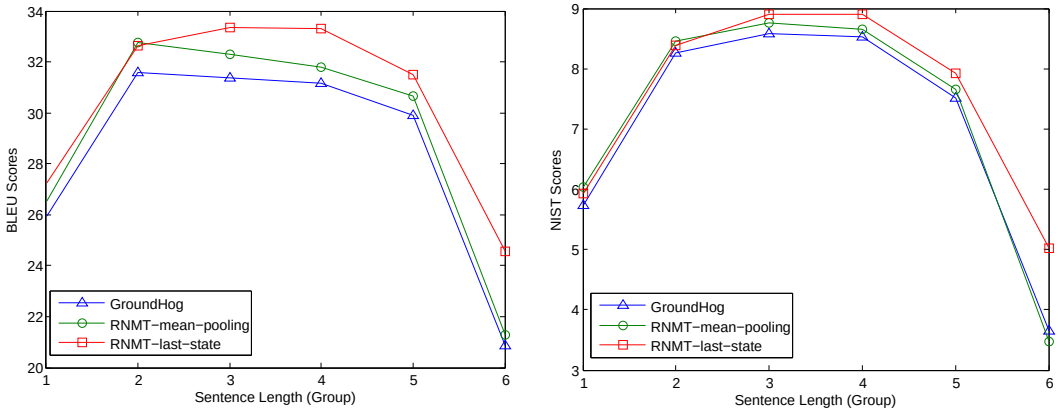


Figure 2: BLEU (left) and NIST (right) scores on different translation groups divided by source sentence length.

our RNMT based on the code of *GroundHog*[4].

All NMT systems are trained on a GeForce GTX TITAN X. In one hour, the *GroundHog* system processes about 1087 batches, while our *RNMT* processes 300 batches. It takes about two weeks to train our model on the training data.

### 4.2 Translation Results

Our first experiment checks whether RNMT can obtain satisfactory performance on the test sets. Table 1 summarizes the experiment results. As the strongest conventional translation system, *Moses* achieves very promising results in terms of both BLUE and NIST (especially it achieves the best results on NIST). This is because *Moses* not only preserves all the words in vocabulary, but also leverages a large monolingual corpus for language model training. In contrast, NMT systems only use a limited vocabulary due to the computation limitation. The performance of *GroundHog* can be regarded as comparable with *Moses*. Replacing the attention mechanism with our contexter, our model achieves significant improvements in terms of both BLEU and NIST. Specifically, our model outperforms the *Moses* 1.31 BLEU points, and the *GroundHog* 1.79 and 0.36 BLEU and NIST points respectively on the whole test set. As the only difference between our *RNMT* and the GroundHog is the *contexter*, these results demonstrate that the recurrent contexter surpasses the attention mechanism.

It is interesting to observe that *RNMT* with *last state* yields better performance than that with *mean-pooling*, with a gain of 1.0 and 0.21 BLEU and NIST points respectively on the whole test set. The reason behind is not completely understood, but can be reflected, to some extent, through the visualization of the *contexter* in the next subsection. Basically, this result suggests that *last state* is more suitable for machine translation.

We argue that our model is able to model long-term dependencies, which is superior to the *GroundHog*. To verify this point, we carried out two addi-

[4]https://github.com/lisa-groundhog/GroundHog

| Dataset | Side | MT02 | MT03 | MT04 | MT05 | MT06 | MT08 |
|---|---|---|---|---|---|---|---|
| *Original* | *Src* | 27.40 | 28.20 | 28.18 | 29.28 | 24.74 | 24.76 |
| | *Ref* | 27.99 | 28.51 | 30.46 | 29.53 | 26.15 | 26.71 |
| *Synthetic* | *Src* | 50.64 | 52.30 | 54.10 | 54.83 | 47.28 | 48.71 |
| | *Ref* | 51.73 | 52.88 | 56.50 | 55.23 | 50.02 | 52.51 |

Table 2: Average source (*Src*) and reference (*Ref*) sentence length for each dataset.

| Metric | System | MT05 | MT02 | MT03 | MT04 | MT06 | MT08 |
|---|---|---|---|---|---|---|---|
| BLEU | *GroundHog* | 18.23 | 22.20 | 20.19 | 21.67 | 19.11 | 13.41 |
| | *RNMT (mean-pooling)* | $19.05^{++}$ | $23.11^{++}$ | $21.17^{++}$ | $23.11^{++}$ | $19.55^{+}$ | $14.16^{++}$ |
| | *RNMT (last state)* | $\mathbf{21.86}^{++}$ | $\mathbf{26.23}^{++}$ | $\mathbf{23.64}^{++}$ | $\mathbf{26.29}^{++}$ | $\mathbf{21.80}^{++}$ | $\mathbf{16.26}^{++}$ |
| NIST | *GroundHog* | 3.23 | 4.25 | 3.35 | 3.80 | 3.69 | 2.23 |
| | *RNMT (mean-pooling)* | 3.27 | 4.34 | 3.51 | $4.00^{+}$ | 3.71 | 2.28 |
| | *RNMT (last state)* | $\mathbf{4.30}^{++}$ | $\mathbf{5.31}^{++}$ | $\mathbf{4.46}^{++}$ | $\mathbf{5.03}^{++}$ | $\mathbf{4.64}^{++}$ | $\mathbf{3.24}^{++}$ |

Table 3: BLEU and NIST scores on the synthetic datasets.

tional experiments.

Our second experiment testifies whether RNMT has better translation quality on long sentences in the original NIST test sets. To this end, we divide our test sets into 6 disjoint groups according to the length of source sentences, each of which has 601, 1918, 1803, 1196, 633 and 455 sentences respectively. Figure 2 illustrates the overall results. We found that as the length of source sentence excels a certain threshold (here from the 5 group to 6 group), the performance of NMT systems drops sharply, around 10 and 4 BLEU and NIST points respectively. This indicates that long sentence translation is a common yet serious challenge for NMT systems. However, compared with the *GroundHog*, our *RNMT* model behaves more robust to long sentences, with a gain of 3.69 and 1.37 BLEU and NIST points respectively on the final group. This further demonstrates that our model indeed deals better with long sentences than the attention-based NMT.

Our third experiment verifies the findings in our second experiment by evaluating our model on a synthetic test data whose source sentence tends to be much longer. We observed that the NIST test sets are organized in different documents. Therefore, we concatenated the original two neighboring sentences in one document as a new sentence for evaluation. In this way, we have the same number of evaluation dataset, but the source and target sentences in the new dataset are much longer. The average sentence length for each dataset is shown in Table 2. Notice that an important property of the new sentence is that it is semantic coherent, thus reasoning inside the sentence is required for correct translation.

Table 3 gives the translation results. We can observe that our *RNMT* especially with *last state* outperforms the *GroudHog* significantly, which gains about 3 and 1 BLEU and NIST points respectively. In some dataset, our model yields above 4 BLEU points. Besides, both of our model achieve improvement over the *GroudHog* no matter which evaluation metric. This result, without any doubt, demonstrates that our model is capable of dealing with long sentence translation.

### 4.3 Contexter Visualization

A common challenge for deep learning is to understand what has happened inside the model. Although our model achieves better performance, we are also curious about the internal mechanism of our contexter. To this end, we provided a method to visualize the *contexter* qualitatively.

Let's revisit the Eq. (8). The GRU consists of two gates: *reset* gate $\mathbf{r}_t$ and *update* gate $\mathbf{z}_t$. Intuitively, the reset gate controls how much information flow from previous state into current state, while the update gate decides how much information should come from current input. That is, the reset gate can reflect the dependencies among the source annotation, while the update gate can qualify how much each source annotation contributes to the word prediction. Therefore, we define two metrics to visualize our *contexter*:

$$\widetilde{r}_t = \frac{1}{d_h}\sum_{k=1}^{d_h}\mathbf{r}_{t,k}$$

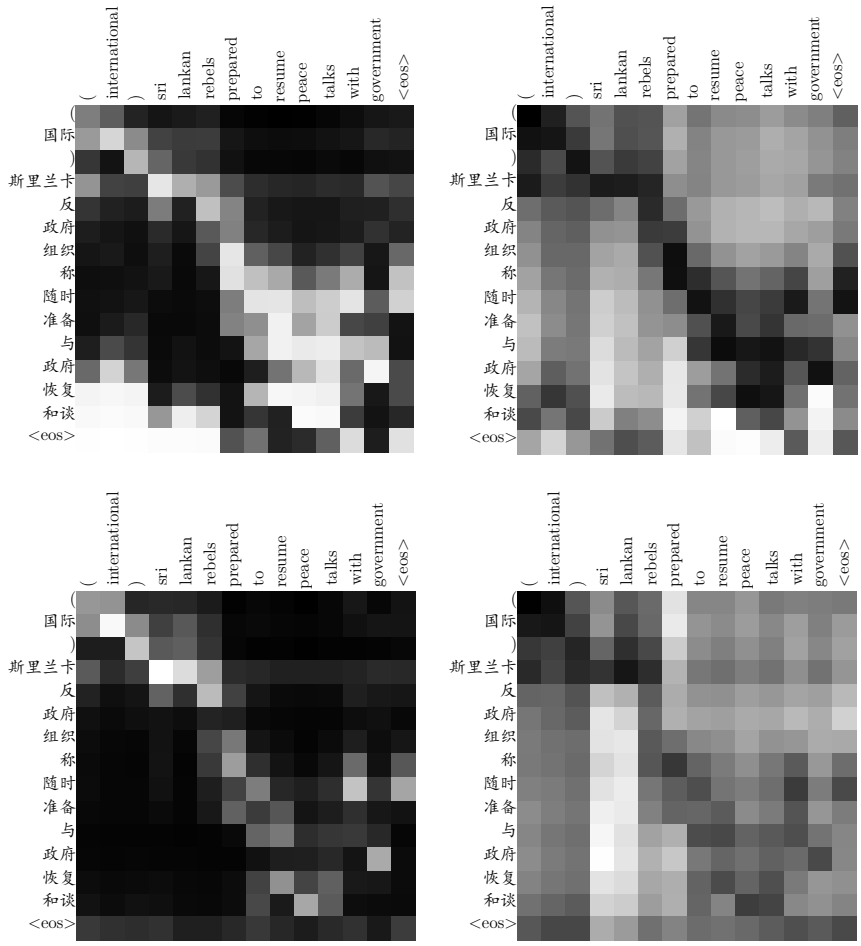


Figure 3: Visualization of the gates in RNMT. The top and bottom row is for *mean-pooling* and *last state* respectively. The left and right column is for *update* and *reset* gate respectively. Notice the white and black blocks for update and reset gate respectively.

$$\widetilde{z}_t = \frac{1}{d_h} \sum_{k=1}^{d_h} \mathbf{z}_{t,k}$$

where $\widetilde{r}_t, \widetilde{z}_t \in (0, 1)$ because of the $\sigma(\cdot)$ in Eq. (8).

The heatmaps of $\widetilde{z}$ and $\widetilde{r}$ in our *RNMT* with *mean-pooling* and *last state* for a same bilingual sentence pair are shown in Figure 3. We find that the update gates differ from the reset gates significantly, which suggests that these two kinds of gates learn different aspects of the source annotations. With regards to the update gate, we find that it looks somewhat like the word alignments, with each block revealing a translation relation from the source to target. From this perspective, the *last state* has a better alignment intuition than the *mean-pooling*.

With respect to the reset gate, we find that it prefers to be opposite to the update gates. For example, the leading diagonal of update gates is white, while that of reset gates is black. Besides, the blocks for aligned word pairs seem to be turn points, which tend to have relatively smaller reset values. This is desirable, since previous hidden states on these points are blocked, while the current hidden states are enabled to flow into the final context vector.

All these indicate that the GRU in our contexter indeed learns something related to translation. Particularly, the contexter implicitly reflects the alignment to source sentence via the reset and update gates. From this perspective, we argue that the ability of automatic word alignment is important for machine translation.

## 5 Related Work

Our work brings together two tightly related line of research: *recurrent neural network* and *neural ma-*

*chine translation*. Although previous work have already testified the effectiveness of the former for the encoder and decoder in the latter, the introduction for the contexter, to the best of our knowledge, has never been investigated before.

*Recurrent Neural Network* takes a variable-length sequence of tokens as its input, and processes each token recursively while maintaining an internal hidden state. Typically, the hidden state is able to encode appealing context information for the history tokens. However, conventional RNNs often suffer from the vanishing and the exploding gradient problems during training (Bengio et al., 1994), which results in difficulties in capturing long-term dependencies. To solve this problem, Hochreiter and Schmidhuber (1997) propose the well-known Long Short-Term Memory (LSTM) which explicitly brings in a memory cell to memorize long-term dependencies together with input, forget and output gates for information controlling. Considering the inefficiency of LSTM in physical computation, Chung et al. (2014) further simplify this architecture, and introduce the gated recurrent unit (GRU) where only a reset and a update gate remains. Both LSTM and GRU have been demonstrated efficient in modeling long-term dependencies through several empirical evaluations (Chung et al., 2014; Bahdanau et al., 2014; Sutskever et al., 2014).

As our goal is to model the long-term dependencies inside source annotations, we borrow the spirit of RNNs and employ the GRU as our contexter in our model.

*Neural Machine Translation*, unlike conventional statistical machine translation, transforms a source sentence directly into its corresponding target reference via a single, large yet trainable neural network. In an early stage, the NMT model follows a *sequence-to-sequence* manner, where a multilayered LSTM encoder encodes the source sentence into a fixed vector, from which a multilayered LSTM decoder generates its target reference word by word until reaching a special end token (Sutskever et al., 2014). A major limitation of this model, however, is that it forces the source-side information within a single vector. To settle this problem, Bahdanau et al. (2014) develop the attention mechanism, which is able to generate target word according to previous target-side hidden state and all the source annotations.

Our work closely relates to the attention-based NMT in the following two aspects: 1) our work borrows the same framework from the attention-based NMT, and 2) our work is also inspired by the basic motivation of the attention mechanism. Despite of its success, current attention-based NMT model still suffers from failures in long sentence translation, e.g. the under-translation problem. Tu et al. (2016) argue that it is the wrong attention that results in these failures, and propose to maintain a coverage vector that encourages NMT to pay more attention to untranslated words and less attention to translated word. In contrast, Cohn et al. (2016) and Feng et al. (2016) explore the "fertility" for the attention mechanism and constrain it with the alignment scores and previous attentional context respectively. Cheng et al. (2015) propose to encourage the agreement on word alignment matrices produced by the source-to-target and target-to-source bidirectional translation models. Nevertheless, what they mainly focus on is the word alignment, rather than the drawback of attention in modeling dependencies among source annotations. We, instead, make up this drawback by replacing the attention mechanism with a more powerful neural network, i.e. GRU.

## 6 Conclusion and Future Work

In this paper, we have presented a recurrent neural machine translation model that replaces the attention mechanism with a gated recurrent unit (GRU). In this way, our model is able to not only leverage the varying-length source annotations, but also handle long-term dependencies among these annotations. Experiment results demonstrate the ability of our model in translating long sentences, and further analysis of the GRU reveals that our model implicitly captures the word alignment.

In the future, we would like to explore two interesting directions: 1) Despite of the effectiveness of our model, the GRU still suffers from computation inefficiency. We will try to optimize its speed. 2) We will explore other neural architectures as the contexter, such as deep neural network, convolutional neural network, bidirectional recurrent neural networks, etc.

## References


[Bahdanau et al.2014] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. In *Proc. of ICLR*.

[Bengio et al.1994] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. 1994. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5:157–166.

[Cheng et al.2015] Yong Cheng, Shiqi Shen, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2015. Agreement-based joint training for bidirectional attention-based neural machine translation. *CoRR*, abs/1512.04650.

[Chung et al.2014] Junyoung Chung, Çaglar Gülçehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*.

[Cohn et al.2016] Trevor Cohn, Cong Duy Vu Hoang, Ekaterina Vymolova, Kaisheng Yao, Chris Dyer, and Gholamreza Haffari. 2016. Incorporating structural alignment biases into an attentional neural translation model. *CoRR*, abs/1601.01085.

[Doddington2002] George Doddington. 2002. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *Proc. of HLT*, pages 138–145.

[Feng et al.2016] Shi Feng, Shujie Liu, Mu Li, and Ming Zhou. 2016. Implicit distortion and fertility models for attention-based encoder-decoder NMT model. *CoRR*, abs/1601.03317.

[Hochreiter and Schmidhuber1997] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9:1735–1780.

[Jean et al.2015] Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation. In *Proc. of ACL-IJCNLP*, pages 1–10, July.

[Koehn et al.2007] Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. of ACL*, pages 177–180.

[Koehn2004] Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proc. of EMNLP*.

[Luong et al.2015a] Thang Luong, Hieu Pham, and Christopher D. Manning. 2015a. Effective approaches to attention-based neural machine translation. In *Proc. of EMNLP*, pages 1412–1421, September.

[Luong et al.2015b] Thang Luong, Ilya Sutskever, Quoc Le, Oriol Vinyals, and Wojciech Zaremba. 2015b. Addressing the rare word problem in neural machine translation. In *Proc. of ACL-IJCNLP*, pages 11–19, July.

[Papineni et al.2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proc. of ACL*, pages 311–318.

[Shen et al.2015] Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2015. Minimum risk training for neural machine translation. *CoRR*, abs/1512.02433.

[Sutskever et al.2014] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. *CoRR*, abs/1409.3215.

[Tu et al.2016] Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Coverage-based neural machine translation. *CoRR*, abs/1601.04811.

[Zhang et al.2016] Biao Zhang, Deyi Xiong, and Jinsong Su. 2016. Variational Neural Machine Translation. *ArXiv e-prints*, May.